

Refine Search

Search Results -

Terms	Documents
L63 and (table with recovery or table near recovery or table adj recovery)	17

Database:

US Pre-Grant Publication Full-Text Database
 US Patents Full-Text Database
 US OCR Full-Text Database
 EPO Abstracts Database
 JPO Abstracts Database
 Derwent World Patents Index
 IBM Technical Disclosure Bulletins

Search:

Refine Search
Recall Text
Clear
Interrupt

Search History

DATE: Wednesday, November 23, 2005 [Printable Copy](#) [Create Case](#)
Set
Name **Query**
 side by
 side

Hit
Count **Set**
Name
 result set

DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR
17 [L64](#)
[L64](#) L63 and (table with recovery or table near recovery or table adj recovery)
129 [L63](#)
[L63](#) L62 and tablespace
6481 [L62](#)
[L62](#) db2

DB=USPT; PLUR=YES; OP=OR
[L61](#) '5222235'.pn.
1 [L61](#)
[L60](#) '5222235'.pn.
1 [L60](#)
[L59](#) '5396623'.pn.
1 [L59](#)
[L58](#) '5396623'.pn.
1 [L58](#)
[L57](#) '5408654'.pn.
1 [L57](#)
[L56](#) '5408654'.pn.
1 [L56](#)
[L55](#) '5710916'.pn.
1 [L55](#)
[L54](#) '5758357'.pn.
1 [L54](#)
[L53](#) '5812849'.pn.
1 [L53](#)
[L52](#) '5937415'.pn.
1 [L52](#)
[L51](#) '5937415'.pn.
1 [L51](#)

<u>L50</u>	'5613106'.pn.	1	<u>L50</u>
<u>L49</u>	'5613106'.pn.	1	<u>L49</u>
<u>L48</u>	'5668986'.pn.	1	<u>L48</u>
<u>L47</u>	'5668986'.pn.	1	<u>L47</u>
<u>L46</u>	'5799141'.pn.	1	<u>L46</u>
<u>L45</u>	'5666530'.pn.	1	<u>L45</u>
<u>L44</u>	'5596706'.pn.	1	<u>L44</u>
<u>L43</u>	'5596706'.pn.	1	<u>L43</u>
<i>DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR</i>			
<u>L42</u>	"parker, christopher".in.	32	<u>L42</u>
<u>L41</u>	5758357.pn.	2	<u>L41</u>
<u>L40</u>	4961134.pn.	2	<u>L40</u>
<u>L39</u>	6119128.pn.	2	<u>L39</u>
<u>L38</u>	6189010.pn.	2	<u>L38</u>
<i>DB=USPT; PLUR=YES; OP=OR</i>			
<u>L37</u>	(5121493 4890226 4679139 5257362 5204958 5274805 5222235 5034914)![PN]	8	<u>L37</u>
<i>DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR</i>			
<u>L36</u>	('5887274' '5758357')[ABPN1,NRPN,PN,TBAN,WKU]	4	<u>L36</u>
<u>L35</u>	L34 and 707/202	4	<u>L35</u>
<u>L34</u>	db2 and tablespace.ab.ti.	489	<u>L34</u>
<i>DB=USPT; PLUR=YES; OP=OR</i>			
<u>L33</u>	'5758357'.pn.	1	<u>L33</u>
<u>L32</u>	(5121493 4890226 4679139 5257362 5204958 5274805 5455945 5414834 5222235 5034914)![PN]	10	<u>L32</u>
<u>L31</u>	('6026412' '5758357')[PN]	2	<u>L31</u>
<u>L30</u>	'5758357'.pn.	1	<u>L30</u>
<u>L29</u>	'5758357'.pn.	1	<u>L29</u>
<u>L28</u>	(5455945 5414834)![PN]	2	<u>L28</u>
<i>DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR</i>			
<u>L27</u>	('6026412')[ABPN1,NRPN,PN,TBAN,WKU]	2	<u>L27</u>
<i>DB=USPT; PLUR=YES; OP=OR</i>			
<u>L26</u>	(5615329 5903898 5640561 5530855 5615364 5729735 5907848 5596706 6070174 6065108 5170480 5325522 6178427 5974563 5924096 5684990 5666530 5455947 5548750)![PN]	19	<u>L26</u>
<i>DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR</i>			
<u>L25</u>	('6289357')[ABPN1,NRPN,PN,TBAN,WKU]	2	<u>L25</u>
<i>DB=USPT; PLUR=YES; OP=OR</i>			
<u>L24</u>	'5907848'.pn.	1	<u>L24</u>
<u>L23</u>	'5907848'.pn.	1	<u>L23</u>
<u>L22</u>	'5924096'.pn.	1	<u>L22</u>
<u>L21</u>	'5924096'.pn.	1	<u>L21</u>
<u>L20</u>	'5974563'.pn.	1	<u>L20</u>
<u>L19</u>	'6065108'.pn.	1	<u>L19</u>
<u>L18</u>	'6070174'.pn.	1	<u>L18</u>
<u>L17</u>	'6065108'.pn.	1	<u>L17</u>

L16 '6065108'.pn. 1 L16
DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR

L15 L14 and db2 near tablespace 1 L15
L14 L13 and (backup or back adj up or back with up or back near up) near (copy or copies or replicat or replicate) 190 L14

L13 L12 and (table with space or table adj space or table-space or tablespace or table near space) 9008 L13
L12 (database or data adj bas or data with base) 952722 L12

L11 707.clas. 30557 L11
L10 707/204 2417 L10
L9 707/203 2862 L9
L8 707/202 2093 L8
L7 707/201 2764 L7
L6 707/200 4115 L6
L5 707/7 1680 L5
L4 707/102 6597 L4
L3 707/101 4205 L3
L2 707/100 6897 L2
L1 707/1 7200 L1

END OF SEARCH HISTORY

[First Hit](#) [Fwd Refs](#)[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)
 [Generate Collection](#) [Print](#)

L14: Entry 120 of 190

File: USPT

Sep 11, 2001

US-PAT-NO: 6289357

DOCUMENT-IDENTIFIER: US 6289357 B1

TITLE: Method of automatically synchronizing mirrored database objects

DATE-ISSUED: September 11, 2001

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Parker; Christopher F.	Oswego	IL		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
Platinum Technology IP, Inc.	Oakbrook Terrace	IL			02

APPL-NO: 09/066183 [PALM]

DATE FILED: April 24, 1998

INT-CL: [07] G06 F 17/30

US-CL-ISSUED: 707/202; 203/204, 203/205, 203/206, 203/201, 203/200

US-CL-CURRENT: 707/202; 707/200, 707/201, 707/203, 707/204, 707/205, 707/206

FIELD-OF-SEARCH: 707/8, 707/10, 707/100, 707/200, 707/201, 707/202, 707/203, 707/204-205, 707/104, 707/206, 714/5, 714/6

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

 [Search Selected](#) [Search All](#) [Clear](#)

PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
<u>5170480</u>	December 1992	Mohan et al.	707/201
<u>5325522</u>	June 1994	Vaughn	707/1
<u>5455947</u>	October 1995	Suzuki et al.	707/204
<u>5530855</u>	June 1996	Satoh et al.	707/201
<u>5546750</u>	August 1996	Larsson et al.	707/204
<u>5596706</u>	January 1997	Shimazaki et al.	714/6
<u>5615329</u>	March 1997	Kern et al.	714/6
<u>5615364</u>	March 1997	Marks	707/202
<u>5640561</u>	June 1997	Satoh et al.	707/202
<u>5666530</u>	September 1997	Clark et al.	707/201

<input type="checkbox"/>	<u>5684990</u>	November 1997	Boothby	707/203
<input type="checkbox"/>	<u>5729735</u>	March 1998	Meyering	707/10
<input type="checkbox"/>	<u>5903898</u>	May 1999	Cohen et al.	704/204
<input type="checkbox"/>	<u>5907848</u>	May 1999	Zaiken et al.	707/202
<input type="checkbox"/>	<u>5924096</u>	July 1999	Draper et al.	707/10
<input type="checkbox"/>	<u>5974563</u>	October 1999	Beeler, Jr.	714/5
<input type="checkbox"/>	<u>6065108</u>	May 2000	Beier et al.	707/202
<input type="checkbox"/>	<u>6070174</u>	May 2000	Starek et al.	707/206
<input type="checkbox"/>	<u>6178427</u>	January 2001	Parker	707/202

FOREIGN PATENT DOCUMENTS

FOREIGN-PAT-NO	PUBN-DATE	COUNTRY	CLASS
403043789A	February 1991	JP	
403256143A	November 1991	JP	
404352061A	December 1992	JP	

ART-UNIT: 211

PRIMARY-EXAMINER: Black; Thomas

ASSISTANT-EXAMINER: Mizrahi; Diane D.

ATTY-AGENT-FIRM: Baker & McKenzie

ABSTRACT:

A method of automatically synchronizing a database with a backup database. The method automatically detects whether the backup database has fallen out of sync with the primary database which it mirrors. In response, a resynchronization process is automatically initiated in which the starting point for reading a redo log of the primary database is automatically determined and in which the backup database is automatically synchronized with the primary dataset.

21 Claims, 8 Drawing figures

[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)

First Hit Fwd Refs

[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)
 [Generate Collection](#) [Print](#)

L26: Entry 1 of 19

File: USPT

Jan 23, 2001

US-PAT-NO: 6178427

DOCUMENT-IDENTIFIER: US 6178427 B1

TITLE: Method of mirroring log datasets using both log file data and live log data including gaps between the two data logs

DATE-ISSUED: January 23, 2001

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Parker; Christopher F.	Oswego	IL		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
Platinum Technology IP, Inc.	Oakbrook Terrace	IL			02

APPL-NO: 09/074080 [PALM]

DATE FILED: May 7, 1998

INT-CL: [07] G06 F 12/00

US-CL-ISSUED: 707/202, 707/203, 707/204

US-CL-CURRENT: 707/202; 707/203, 707/204

FIELD-OF-SEARCH: 707/202, 707/204, 707/203

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

 [Search Selected](#) [Search All](#) [Clear](#)

PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
<u>5544359</u>	August 1996	Tada et al.	395/600
<u>5613106</u>	March 1997	Thurman et al.	395/620
<u>5668986</u>	September 1997	Nilsen et al.	395/610
<u>5742792</u>	April 1998	Yanai et al.	395/489
<u>5799141</u>	August 1998	Galipeau et al.	395/182
<u>5901327</u>	May 1999	Ofek	395/825
<u>5903898</u>	November 1999	Cohen et al.	707/204
<u>6044444</u>	March 2000	Ofek	711/162

ART-UNIT: 271

PRIMARY-EXAMINER: Black; Thomas G.

ASSISTANT-EXAMINER: Do; Thuy

ATTY-AGENT-FIRM: Baker & McKenzie

ABSTRACT:

A method of mirroring log data in real time. The method can be used to mirror all log data from a source database or only that log data pertaining to a subset of objects in the source database. The log data is written to one or more datasets which are created and stored in a target database system which is typically remotely located. The log datasets in the target system can be used for disaster recovery of the source system. The method monitors log data from the source system and automatically updates target log datasets accordingly.

6 Claims, 6 Drawing figures

[Previous Doc](#)

[Next Doc](#)

[Go to Doc#](#)

PORTAL

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library The Guide

USPTO + "transactions on database systems"

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used transactions on database systems

Found 3,958 of 167,655

Sort results by relevance Save results to a Binder
 Display results expanded form Search Tips
 Open results in a new window

Try an Advanced Search
 Try this search in The ACM Guide

Results 1 - 20 of 200

Result page: 1 2 3 4 5 6 7 8 9 10 next

Best 200 shown

Relevance scale

1 Modeling the storage architectures of commercial database systems

D. S. Batory

December 1985 **ACM Transactions on Database Systems (TODS)**, Volume 10 Issue 4

Publisher: ACM Press

Full text available:  pdf(4.46 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Modeling the storage structures of a DBMS is a prerequisite to understanding and optimizing database performance. Previously, such modeling was very difficult because the fundamental role of conceptual-to-internal mappings in DBMS implementations went unrecognized. In this paper we present a model of physical databases, called the transformation model, that makes conceptual-to-internal mappings explicit. By exposing such mappings, we show that it is possible to model the storage ...

2 Broadcast protocols to support efficient retrieval from databases by mobile users

Anindya Datta, Debra E. VanderMeer, Aslihan Celik, Vijay Kumar

March 1999 **ACM Transactions on Database Systems (TODS)**, Volume 24 Issue 1

Publisher: ACM Press

Full text available:  pdf(638.48 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Mobile computing has the potential for managing information globally. Data management issues in mobile computing have received some attention in recent times, and the design of adaptive broadcast protocols has been posed as an important problem. Such protocols are employed by database servers to decide on the content of broadcasts dynamically, in response to client mobility and demand patterns. In this paper we design such protocols and also propose efficient retrieval s ...

Keywords: adaptive broadcast protocols, client-server computing, energy conservation, mobile databases

3 An algebraic approach to static analysis of active database rules

Elena Baralis, Jennifer Widom

September 2000 **ACM Transactions on Database Systems (TODS)**, Volume 25 Issue 3

Publisher: ACM Press

Full text available:  pdf(391.93 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Rules in active database systems can be very difficult to program due to the unstructured and unpredictable nature of rule processing. We provide static analysis techniques for predicting whether a given rule set is guaranteed to terminate and whether rule execution

is confluent (guaranteed to have a unique final state). Our methods are based on previous techniques for analyzing rules in active database systems. We improve considerably on the previous techniques by providing analysis criter ...

Keywords: active database systems, confluence, database rule processing, database trigger processing, termination

4 ARIES: a transaction recovery method supporting fine-granularity locking and partial

rollbacks using write-ahead logging

C. Mohan, Don Haderle, Bruce Lindsay, Hamid Pirahesh, Peter Schwarz

March 1992 **ACM Transactions on Database Systems (TODS)**, Volume 17 Issue 1

Publisher: ACM Press

Full text available:  pdf(5.23 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

DB2TM, IMS, and TandemTM systems. ARIES is applicable not only to database management systems but also to persistent object-oriented languages, recoverable file systems and transaction-based operating systems. ARIES has been implemented, to varying degrees, in IBM's OS/2TM Extended Edition Database Manager, DB2, Workstation Data Save Facility/VM, Starburst and QuickSilver, and in the University of Wisconsin's EXODUS and Gamma d ...

Keywords: buffer management, latching, locking, space management, write-ahead logging

5 Index-driven similarity search in metric spaces

Gisli R. Hjaltason, Hanan Samet

December 2003 **ACM Transactions on Database Systems (TODS)**, Volume 28 Issue 4

Publisher: ACM Press

Full text available:  pdf(650.64 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Similarity search is a very important operation in multimedia databases and other database applications involving complex objects, and involves finding objects in a data set S similar to a query object q , based on some similarity measure. In this article, we focus on methods for similarity search that make the general assumption that similarity is represented with a distance metric d . Existing methods for handling similarity search in this setting typically fall into one of ...

Keywords: Hierarchical metric data structures, distance-based indexing, nearest neighbor queries, range queries, ranking, similarity searching

6 Supporting valid-time indeterminacy

Curtis E. Dyer, Richard Thomas Snodgrass

March 1998 **ACM Transactions on Database Systems (TODS)**, Volume 23 Issue 1

Publisher: ACM Press

Full text available:  pdf(516.09 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In valid-time indeterminacy it is known that an event stored in a database did in fact occur, but it is not known exactly when. In this paper we extend the SQL data model and query language to support valid-time indeterminacy. We represent the occurrence time of an event with a set of possible instants, delimiting when the event might have occurred, and a probability distribution over that set. We also describe query language constructs to retrieve informat ...

Keywords: SQL, TSQL2, incomplete information, indeterminacy, probabilistic information, temporal database, valid-time database

7 An adaptive data replication algorithm Ouri Wolfson, Sushil Jajodia, Yixiu HuangJune 1997 **ACM Transactions on Database Systems (TODS)**, Volume 22 Issue 2**Publisher:** ACM PressFull text available:  pdf(911.08 KB)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This article addresses the performance of distributed database systems. Specifically, we present an algorithm for dynamic replication of an object in distributed systems. The algorithm is adaptive in the sense that it changes the replication scheme of the object i.e., the set of processors at which the object is replicated) as changes occur in the read-write pattern of the object (i.e., the number of reads and writes issued by each processor). The algorithm continuously moves the replication ...

Keywords: computer networks, dynamic data allocation, file allocation, replicated data**8 Transitive closure algorithms based on graph traversal** Yannis Ioannidis, Raghu Ramakrishnan, Linda WingerSeptember 1993 **ACM Transactions on Database Systems (TODS)**, Volume 18 Issue 3**Publisher:** ACM PressFull text available:  pdf(4.34 MB)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Several graph-based algorithms have been proposed in the literature to compute the transitive closure of a directed graph. We develop two new algorithms (Basic_TC and Gobal_DFTC) and compare the performance of their implementations in a disk-based environment with a well-known graph-based algorithm proposed by Schmitz. Our algorithms use depth-first search to traverse a graph and a technique called marking to avoid processing some of the arcs in the graph. They compute the ...

Keywords: depth-first search, node reachability, path computations, transitive closure**9 Distributed query evaluation on semistructured data** Dan SuciuMarch 2002 **ACM Transactions on Database Systems (TODS)**, Volume 27 Issue 1**Publisher:** ACM PressFull text available:  pdf(689.88 KB)Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Semistructured data is modeled as a rooted, labeled graph. The simplest kinds of queries on such data are those which traverse paths described by regular path expressions. More complex queries combine several regular path expressions, with complex data restructuring, and with sub-queries. This article addresses the problem of efficient query evaluation on distributed, semistructured databases. In our setting, the nodes of the database are distributed over a fixed number of sites, and the ...

Keywords: Distributed evaluation, nested queries, parallel complexity, regular expressions, semistructured data**10 Optimizing object queries using an effective calculus** Leonidas Fegaras, David MaierDecember 2000 **ACM Transactions on Database Systems (TODS)**, Volume 25 Issue 4**Publisher:** ACM PressFull text available:  pdf(641.65 KB)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Object-oriented databases (OODBs) provide powerful data abstractions and modeling facilities, but they generally lack a suitable framework for query processing and optimization. The development of an effective query optimizer is one of the key factors for OODB systems to successfully compete with relational systems, as well as to meet the performance requirements of many nontraditional applications. We propose an effective framework with a solid theoretical basis for optimizing OODB query ...

Keywords: nested relations, object-oriented databases, query decorrelation, query optimization

11 Adaptive, fine-grained sharing in a client-server OODBMS: a callback-based approach

 Markos Zaharioudakis, Michael J. Carey, Michael J. Franklin
December 1997 **ACM Transactions on Database Systems (TODS)**, Volume 22 Issue 4

Publisher: ACM Press

Full text available:  pdf(441.80 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

For reasons of simplicity and communication efficiency, a number of existing object-oriented database management systems are based on page server architectures; data pages are their minimum unit of transfer and client caching. Despite their efficiency, page servers are often criticized as being too restrictive when it comes to concurrency, as existing systems use pages as the minimum locking unit as well. In this paper we show how to support object-level locking in a page-server context. Sev ...

Keywords: cache coherency, cache consistency, client-server databased, fine-grained sharing, object-oriented databases, performance analysis

12 Object normal forms and dependency constraints for object-oriented schemata

 Zahir Tari, John Stokes, Stefano Spaccapietra
December 1997 **ACM Transactions on Database Systems (TODS)**, Volume 22 Issue 4

Publisher: ACM Press

Full text available:  pdf(439.37 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

We address the development of a normalization theory for object-oriented data models that have common features to support objects. We first provide an extension of functional dependencies to cope with the richer semantics of relationships between objects, called path dependency, local dependency, and global dependency constraints. Using these dependency constraints, we provide normal forms for object-oriented data models based on the notions of user ...

Keywords: data model, functional and multivalued dependencies, normal forms, object-oriented paradigm

13 Automatic generation of production rules for integrity maintenance

 Stefano Ceri, Piero Fraternali, Stefano Paraboschi, Letizia Tanca
September 1994 **ACM Transactions on Database Systems (TODS)**, Volume 19 Issue 3

Publisher: ACM Press

Full text available:  pdf(3.42 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

In this article we present an approach to integrity maintenance, consisting of automatically generating production rules for integrity enforcement. Constraints are expressed as particular formulas of Domain Relational Calculus; they are automatically translated into a set of repair actions, encoded as production rules of an active database system. Production rules may be redundant (they enforce the same constraint in different ways) and conflicting (because repairing one constraint may caus ...

Keywords: automatic generation of production rules

14 SilkRoute: A framework for publishing relational data in XML

 Mary Fernández, Yana Kadiyska, Dan Suciu, Atsuyuki Morishima, Wang-Chiew Tan
December 2002 **ACM Transactions on Database Systems (TODS)**, Volume 27 Issue 4

Publisher: ACM Press

Full text available:  pdf(687.91 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

XML is the "lingua franca" for data exchange between interenterprise applications. In this work, we describe SilkRoute, a framework for publishing relational data in XML. In SilkRoute, relational data is published in three steps: the relational tables are presented to the database administrator in a canonical XML view; the database administrator defines in the XQuery query language a public, virtual XML view over the canonical XML view; and an application formulates an XQuery query over the publ ...

Keywords: XML, XML storage systems, XQuery

15 Atomicity and isolation for transactional processes

 Heiko Schuld, Gustavo Alonso, Catriel Beeri, Hans-Jörg Schek
March 2002 **ACM Transactions on Database Systems (TODS)**, Volume 27 Issue 1

Publisher: ACM Press

Full text available:  pdf(1.22 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Processes are increasingly being used to make complex application logic explicit. Programming using processes has significant advantages but it poses a difficult problem from the system point of view in that the interactions between processes cannot be controlled using conventional techniques. In terms of recovery, the steps of a process are different from operations within a transaction. Each one has its own termination semantics and there are dependencies among the different steps. Regarding c ...

Keywords: Advanced transaction models, business process management, electronic commerce, execution guarantees, locking, processes, semantically rich transactions, transactional workflows, unified theory of concurrency control and recovery

16 Probabilistic temporal databases. I: algebra

 Alex Dekhtyar, Robert Ross, V. S. Subrahmanian
March 2001 **ACM Transactions on Database Systems (TODS)**, Volume 26 Issue 1

Publisher: ACM Press

Full text available:  pdf(878.03 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Dyreson and Snodgrass have drawn attention to the fact that, in many temporal database applications, there is often uncertainty about the start time of events, the end time of events, and the duration of events. When the granularity of time is small (e.g., milliseconds), a statement such as "Packet p was shipped sometime during the first 5 days of January, 1998" leads to a massive amount of uncertainty ($5 \times 24 \times 60 \times 60 \times 1000$) possibilities. A ...

17 A structured approach for the definition of the semantics of active databases

 Piero Fraternali, Letizia Tanca
December 1995 **ACM Transactions on Database Systems (TODS)**, Volume 20 Issue 4

Publisher: ACM Press

Full text available:  pdf(4.15 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Active DBMSs couple database technology with rule-based programming to achieve the capability of reaction to database (and possibly external) stimuli, called events. The reactive capabilities of active databases are useful for a wide spectrum of applications, including security, view materialization, integrity checking and enforcement, or heterogeneous database integration, which makes this technology very promising for the near future. An active database system consists of ...

Keywords: active database systems, database rule processing, events, fixpoint semantics, rules, semantics

18 Rule-based optimization and query processing in an extensible geometric database system

Ludger Becker, Ralf Hartmut Güting

June 1992 **ACM Transactions on Database Systems (TODS)**, Volume 17 Issue 2

Publisher: ACM Press

Full text available:  pdf(3.35 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Gral is an extensible database system, based on the formal concept of a many-sorted relational algebra. Many-sorted algebra is used to define any application's query language, its query execution language, and its optimization rules. In this paper we describe Gral's optimization component. It provides (1) a sophisticated rule language—rules are transformations of abstract algebra expressions, (2) a general optimization framework under which more specific optimization algorithms can be ...

Keywords: extensibility, geometric query processing, many-sorted algebra, optimization, relational algebra, rule-based optimization

19 Understanding the global semantics of referential actions using logic rules

Wolfgang May, Bertram Ludäscher

December 2002 **ACM Transactions on Database Systems (TODS)**, Volume 27 Issue 4

Publisher: ACM Press

Full text available:  pdf(640.93 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Referential actions are specialized triggers for automatically maintaining referential integrity in databases. While the *local effects* of referential actions can be grasped easily, it is far from obvious what the *global semantics* of a set of interacting referential actions should be. In particular, when using procedural execution models, ambiguities due to the execution ordering can occur. No *global, declarative* semantics of referential actions has yet been defined. We show t ...

Keywords: Database theory, game theory, logic programming, referential actions, referential integrity, relational databases

20 Temporal statement modifiers

Michael H. Böhlen, Christian S. Jensen, Richard Thomas Snodgrass

December 2000 **ACM Transactions on Database Systems (TODS)**, Volume 25 Issue 4

Publisher: ACM Press

Full text available:  pdf(317.23 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

A wide range of database applications manage time-varying data. Many temporal query languages have been proposed, each one the result of many carefully made yet subtly interacting design decisions. In this article we advocate a different approach to articulating a set of requirements, or desiderata, that directly imply the syntactic structure and core semantics of a temporal extension of an (arbitrary) nontemporal query language. These desiderata facilitate transitioning applications from a ...

Keywords: ATSQL, statement modifiers, temporal databases

Results 1 - 20 of 200

Result page: 1 2 3 4 5 6 7 8 9 10 next

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)